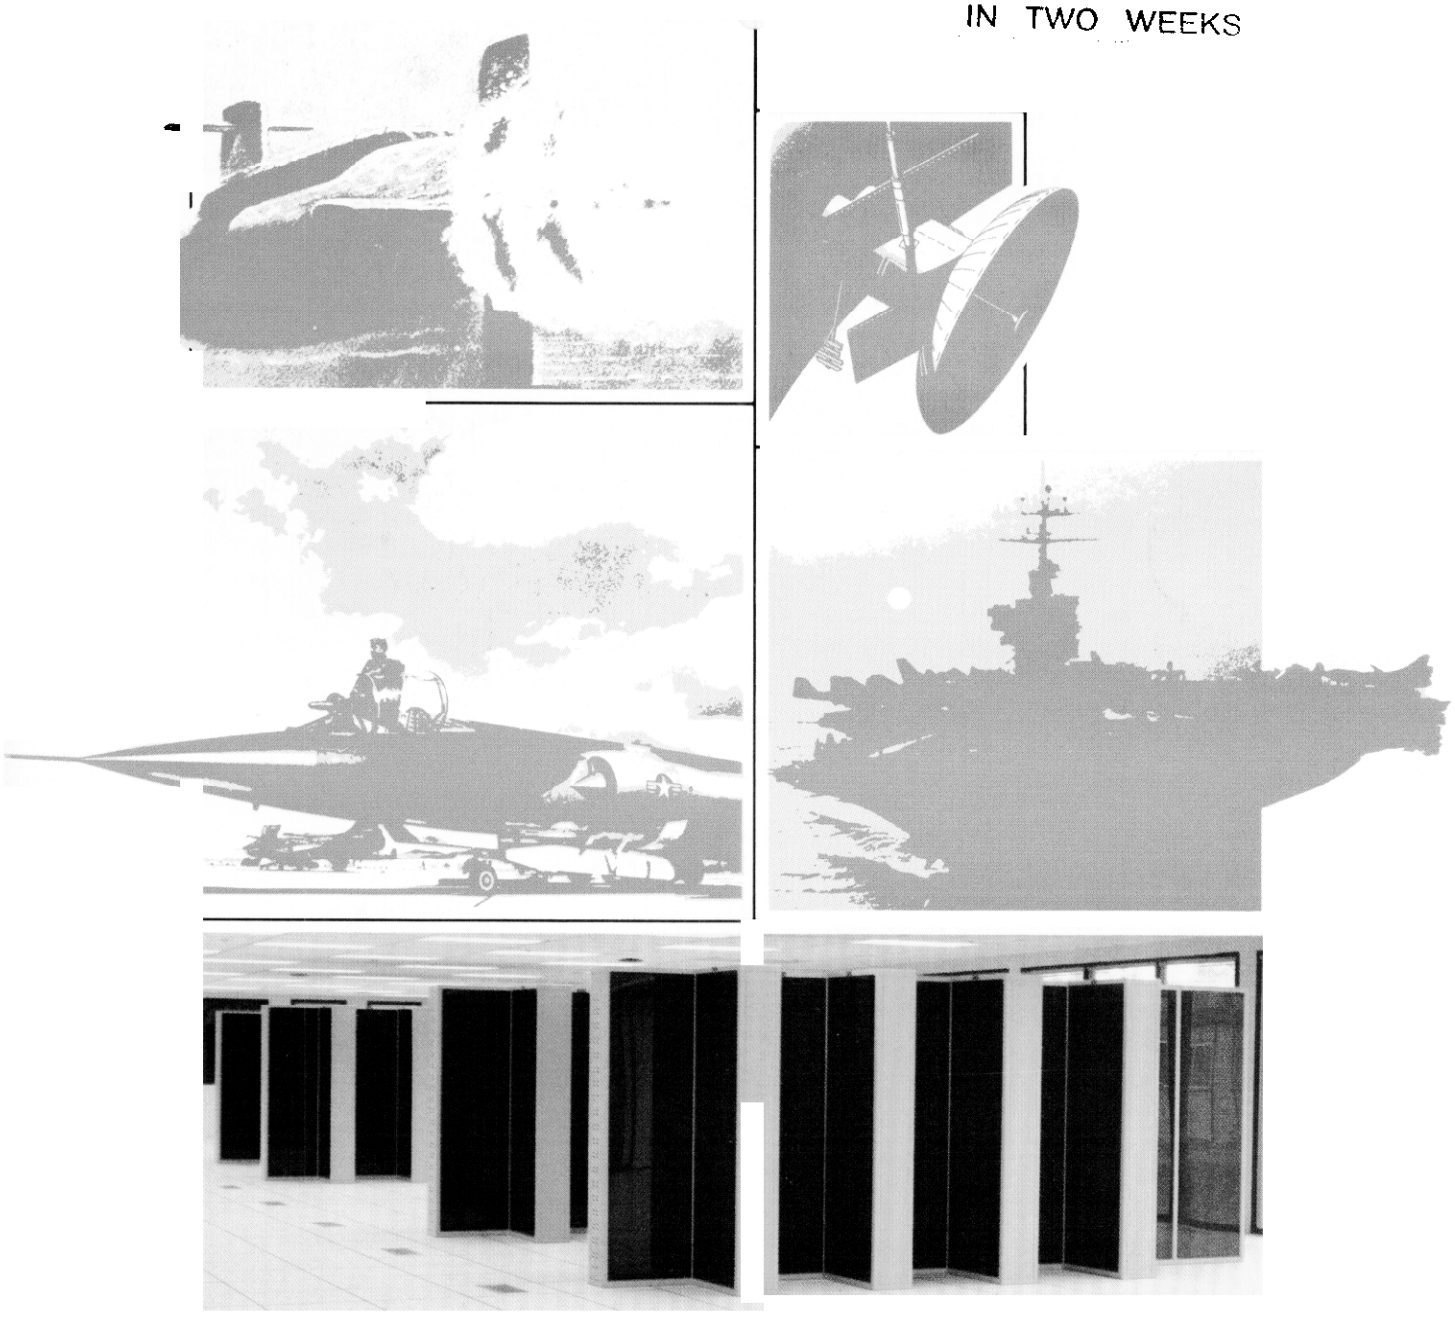

The S-1 Multiprocessor Test and Evaluation Facility

DISCLAIMER

The problem

Better sensors, bigger databases, more complicated weapons systems: all of these demand more computer power. For theater weather forecasting, three-dimensional flow modelling, real-time combat environment simulation, and a host of other applications, current computing systems are just not fast and cheap enough.

The solution

At Lawrence Livermore National Laboratory, NAVELEX is sponsoring a project called "S-1" to develop computer-aided design and manufacturing tools that will advance the state of the art in digital hardware. The first system produced with those tools is almost ready: a unique new multiprocessor that promises not only unprecedented computing power in a single location, but also remarkable cost-effectiveness.

Together with its software, this system is representative of the technology that is expected to be commercially available in the 1990s, technology that a 1981 Defense Science Board study predicted could have a revolutionary impact on warfare.

The plan

We want to measure and evaluate the way this system performs when confronted with the most challenging digital computing problems throughout the Navy. Between now and FY87, NAVMAT is offering to support your use of the S-1 multiprocessor facility. This is your opportunity to measure the impact that advanced computing resources will have on the way you perform missions in the 1990s.

The schedule

The first S-1 system, a Mark IIA uniprocessor, having a scientific computing capability comparable to that of a CRAY-1 plus 128 million bytes of memory, will be available for initial studies in the summer of 1983. It will be available for timesharing under the version 7 Unix operating system in the autumn. The first S-1 multiprocessor, consisting of two uniprocessors and two memories interconnected through a high-bandwidth crossbar switch, will be available in early 1984 under an advanced operating system called Amber. At least four additional Mark IIA uniprocessors and six additional memories will be integrated into the multiprocessor soon afterward.

The Mark IIB

A 5 inch silicon wafer

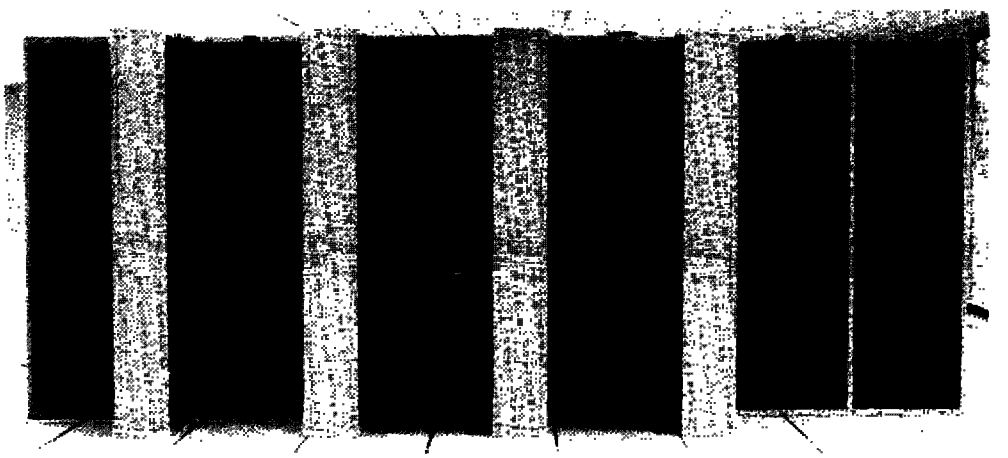

Shrinking circuitry, growing computer power:

The Mark IIA uniprocessor fills a kitchen-sized room; the Mark IIB version (architecturally identical, but faster) will be about the size of a microwave oven; and by 1990, equivalent computer power will be available on a single silicon wafer the size of a saucer.

Side view of the Mark IIA uniprocessor.

Eight independent, loadsharing power supplies are stacked at each end of the cabinet. 72 wire-wrapped panels carry the integrated circuits, with every IC package available for diagnostic probing while the machine is operating. A row of fans at the base of the cabinet pulls chilled air from below the floor and circulates it past the boards.

The uniprocessor

The S-1 Mark IIA uniprocessor executes up to 12.5 million instructions per second. In vector operation, this single processor performs 50 million floating-point operations per second. On particular operations, like the fast fourier transform, it can perform 250 million floating-point operations per second.

But this is a sophisticated computing machine, not just a fast one:

- A 2^{31} byte address space frees programmers from space constraints.
- A 2^{34} byte physical address space together with high-speed data and instruction caches take full advantage of dramatic progress in MOS memory technology.
- Hardware support for segmentation and tracing ease the debugging of large, complex programs.
- Arithmetic hardware pinpoints exception conditions, even allowing users to substitute software "fixup" algorithms for the standard ones.
- Hardware support for procedure calls and addressing reduce the performance penalty for high-level languages.
- Entire algorithms—FFT, vector convolution, matrix multiplication, arctangent, second-order recursive filter, exponential, the quicksort inner loop—are provided as single instructions. Reciprocate and square root instructions are nearly as fast as addition or multiplication.

Uniprocessor software

Unix™ will be available on the Mark IIA from the very start. Originally developed at Bell Laboratories, this timesharing system is now available on a wider variety of computer hardware than any other, and is popular in universities and industry alike. An enormous amount of software has already been designed to run under Unix, and a large number of programmers are already familiar with it.

Later, the Amber operating system will provide a stepping-stone to multiprocessor operation. Amber and Unix share a common ancestor, MIT's pioneering Multics project. But Amber excels in security, real-time capability, and management of large applictions. Through a layered approach, the Amber base system can become a family of systems, addressing various needs: timesharing, real-time, embedded control.

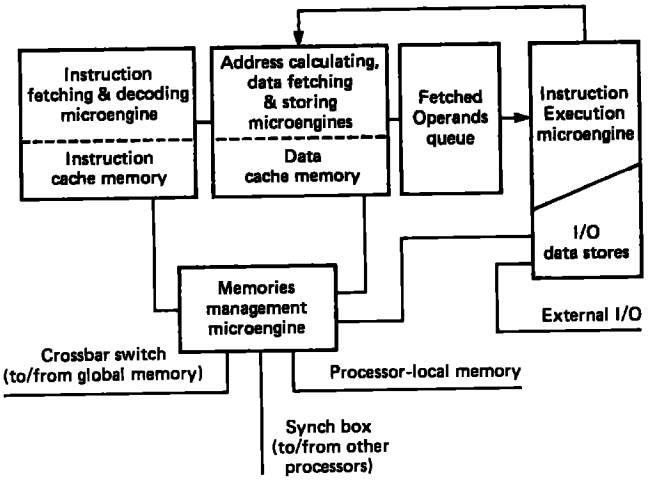

Each of the main functional units of the Mark IIA contains one or more "microengines" which make possible a sophisticated, adaptible architecture. The remarkable input-output rate of the CPU—sustainable at 450 million bits/s—is in turn dwarfed by the 1.8 billion bits/s rate at which data flows between cache memory and the arithmetic units.

The test and evaluation facility will provide a timesharing system, with screen-oriented editors, text processors, and mail systems to enhance programmer productivity.

Language support includes Pascal, C, Fortran and a superset of Common Lisp.

Some 8 ns, 1K-bit ECL memories used in the Mark IIA's microengines.

Uniprocessor Vital Statistics

	Mark IIA	Mark IIB
Physical characteristics		
Footprint, ft^2	25	4
Weight, lb	2500	225
Power consumption, kW	13	10
Components and technology		
IC type	ECL 10K/100K DIP	ECL 100K Flatpak
Arithmetic unit cycle period, ns	40	20
Cooling technology	Forced air	Chilled water/ Forced air
Memory density, Kbits per IC	64	256
Construction technique	Wire wrap	Printed circuit
IC population	27000	20000
Data cache capacity, Kbytes	256	1024
Complex interrupt servicing rate, sec^{-1}	10^5	2×10^5
Opcodes	4096	4096
Maximum signal progessing rate, megaflops	250	512
Instruction timing (expressed in arithmetic unit cycles, for vector pipeline/scalar pipeline/ total latency)	**36-bit data**	**72 bit data**
Move (register to memory)	NA/2/2	NA/2/2
Move (memory to register)	1/4/4	2/4/4
Integer add	1/2/4	2/2/4
Integer multiply	1/2/6	2/2/6
Floating point add	1/2/4	2/2/4
Floating point multiply	1/2/6	2/2/6
Floating point reciprocate	1/2/6	3/14/18
Floating point logarithm	1/8/10	14/46/50
Floating point sine	2/16/20	14/38/42

The multiprocessor

The multiprocessor consists of a number of uniprocessors, each with an associated memory, communicating through a full N^2 crossbar switch. Each uniprocessor has uniform access to all the memory in the multiprocessor; a pair of caches in each CPU hide the latency of main memory, and guarantee coherency when multiple processors share memory.

Multiprocessors like the S-1 present today's best opportunity for dramatically increasing computer power, but they also pose a challenge to users seeking to exploit that increase.

To smooth the transition, the Amber operating system lets the user initially view the system conventionally: at any instant, it runs his application within a particular uniprocessor. Later, to achieve greater speed, the user can employ partitioning tools to analyze his application and, by inserting "hints" in the source code, organize it to run in parallel on any desired number of processors up to the entire system.

Regardless of the number of processors an application requires, the user can ask Amber to give it any desired level of service, from low-priority background operation in a timesharing environment up to exclusive use of the processors.

Next to speed, reliability is the most important advantage of a multiprocessor. In case of processor or memory failure, the system can be rapidly reconfigured to resume operation. I/O processors can be connected redundantly so that failure of a CPU does not isolate I/O devices, and failure of an I/O processor does not isolate a CPU.

The initial multiprocessor configuration at the facility will consist of 6 Mark IIA uniprocessors, 8 memories, and one crossbar switch.

Later, we will add additional memories and Mark IIB versions of the uniprocessor.

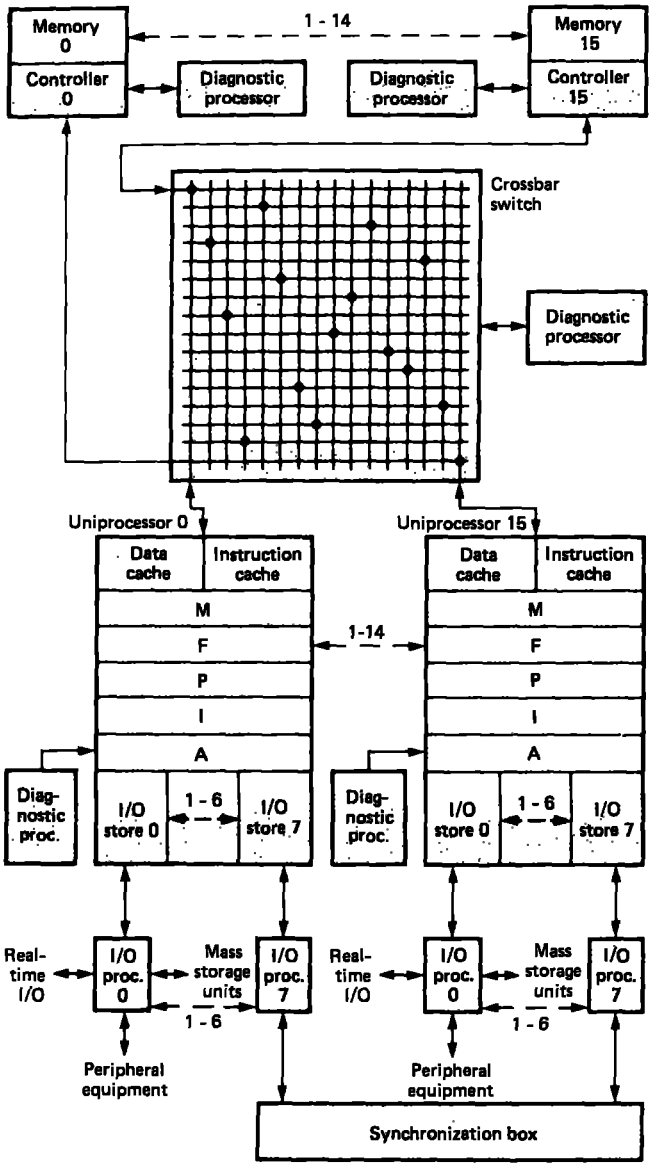

Architecture of an S-1 multiprocessor, showing only two processors, two memories, and one crossbar switch. The initial hardware will allow any configuration from two to sixteen processors.

Sample multiprocessor configurations:

A synthetic-aperture radar configuration

- 2 or more Mark IIA uniprocessors
- 1 crossbar switch/synch box
- 2 or more memories
- 2 or more high-resolution, color graphics displays with hard-copy generation capability
- 2 or more high-bandwidth (IBM 3380 class) disc drives
- 2 or more high-density (6250 bpi) tape drives

Synthetic aperture radar data processing makes extensive use of vector arithmetic operations (mostly addition and multiplication) and fast fourier transforms. A typical SAR data processing configuration would include at least two processors (1 for image formatting and supervising I/O, the others for executing the core SAR processing algorithm) connected to high-performance mass-storage units bringing the raw data into the multiprocessor and carrying synthesized images out of it to display and hard-copy units. Each Mark IIA is estimated to be able to generate a 6000 point by 6000 point SEASAT-A type image from raw data in less than 3 minutes, executing at an average processing rate of about 200 million floating-point operations/s.

A C^2 configuration

- 4 or more Mark IIA uniprocessors
- 2 crossbar switch/synch boxes
- 8 memories
- 4 or more high-resolution, color graphics displays with hard-copy generation capability
- 2 or more high-bandwidth (IBM 3380 class) disc drives
- 2 high-density (6250 bpi) tape drives
- Amber Base Operating System with program development and real time layers, hosting S-1 Lisp environment

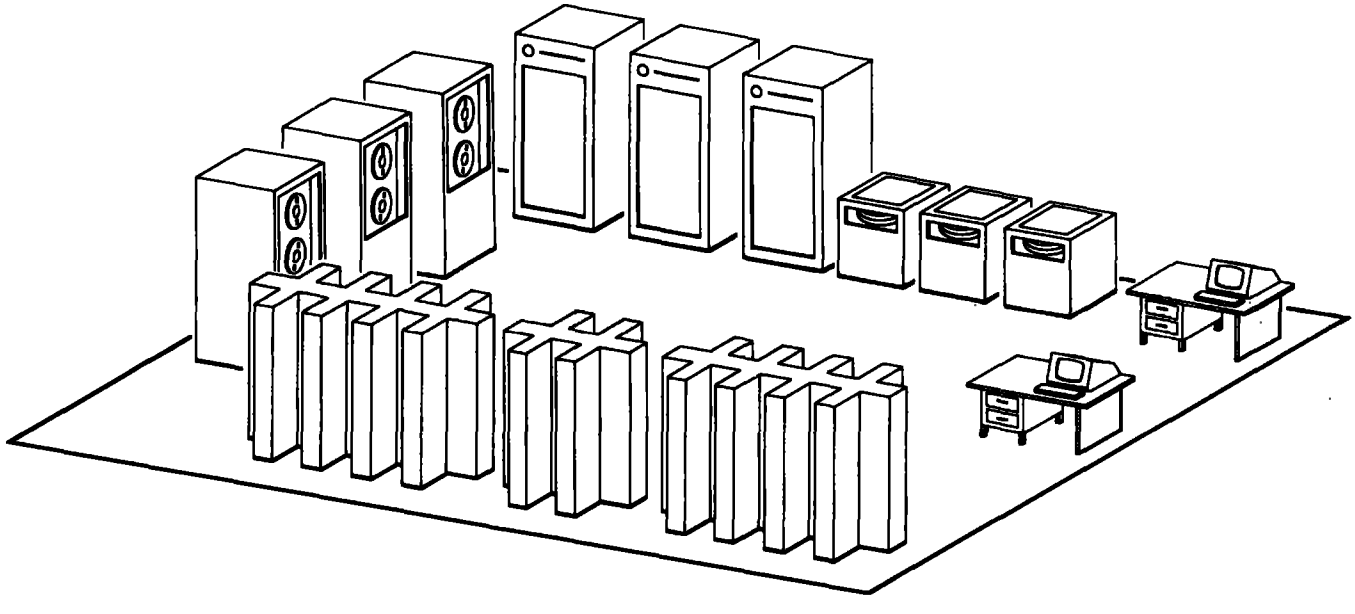

A synthetic-aperture radar configuration

Command/control computing emphasizes continuous system availability, high bandwidth at a user-friendly man-machine interface, high-capacity and high-bandwidth database storage in fully redundant high-speed and disc memory units, and a fully equipped software environment for development, maintenance, and evaluation of C/C programs. Bandwidth between a uniprocessor and memory is 0.5 billion bits per second. Interrupt servicing rate is well in excess of 100,000 per second per uniprocessor. Scientific functions, particularly the trigonometric ones required for target tracking, are implemented as single hardware instructions, and are computed in vector mode at 25 million/s.

A database management configuration

- 2 or more Mark IIA uniprocessors
- 1 or more crossbar switch/synch boxes
- 8 memories
- 2 or more high-resolution, color graphics displays with hard-copy generation capability
- 8 or more high-bandwidth (IBM 3380 class) disc drives
- 2 or more high-density (6250 bpi) tape drives
- Amber Base Operating System with program development layer, hosting S-1 Pascal environment

Database management operations requiring manipulation of very large files are facilitated by the 128 megabytes of SECDED high-speed storage per memory — more capacity than that of all but the largest disc drives — the high bandwidth between each processor and memory, and the even higher data bandwidth internal to each uniprocessor. The inner loop of quicksort is implemented as a single hardware instruction, and is pipelined to provide several times higher sustained sorting rates than on any other computer in existence.

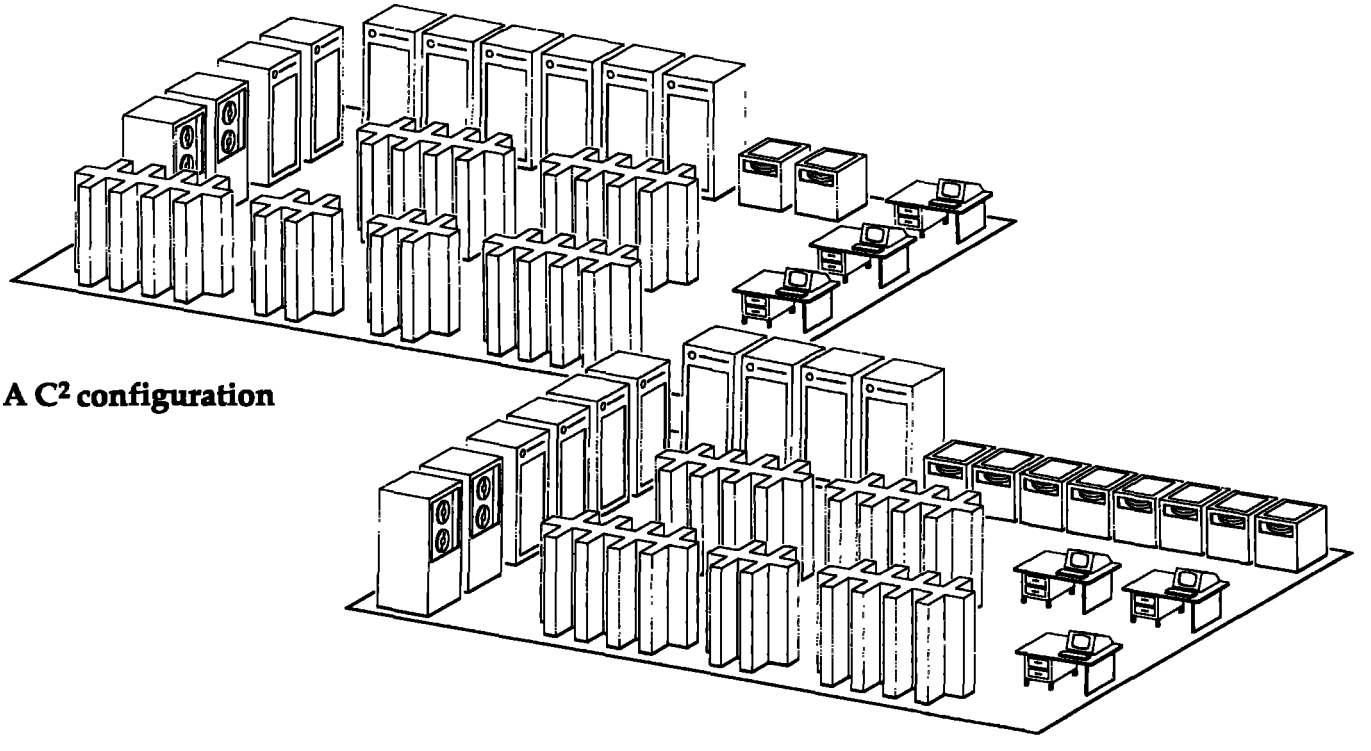

A C^2 configuration

A database management configuration

The offer

We want you to try out this computer system.

Once NAVMAT has approved your request, the system is yours to use throughout the evaluation period, without charge for access time, computing time, or nominal quantities of the storage and output media provided at the facility.

You will be able to access the facility via ARPAnet or MILnet, both of which support classified subnets. Or you can send personnel to Livermore, California for hands-on operation with close-up support by the S-1 Project staff. You must defray your own personnel and travel costs; both the Navy SysComs and ONR have programmed funding for such purposes in FY84-87.

The building housing the S-1 multiprocessor test and evaluation facility at Lawrence Livermore National Laboratory.

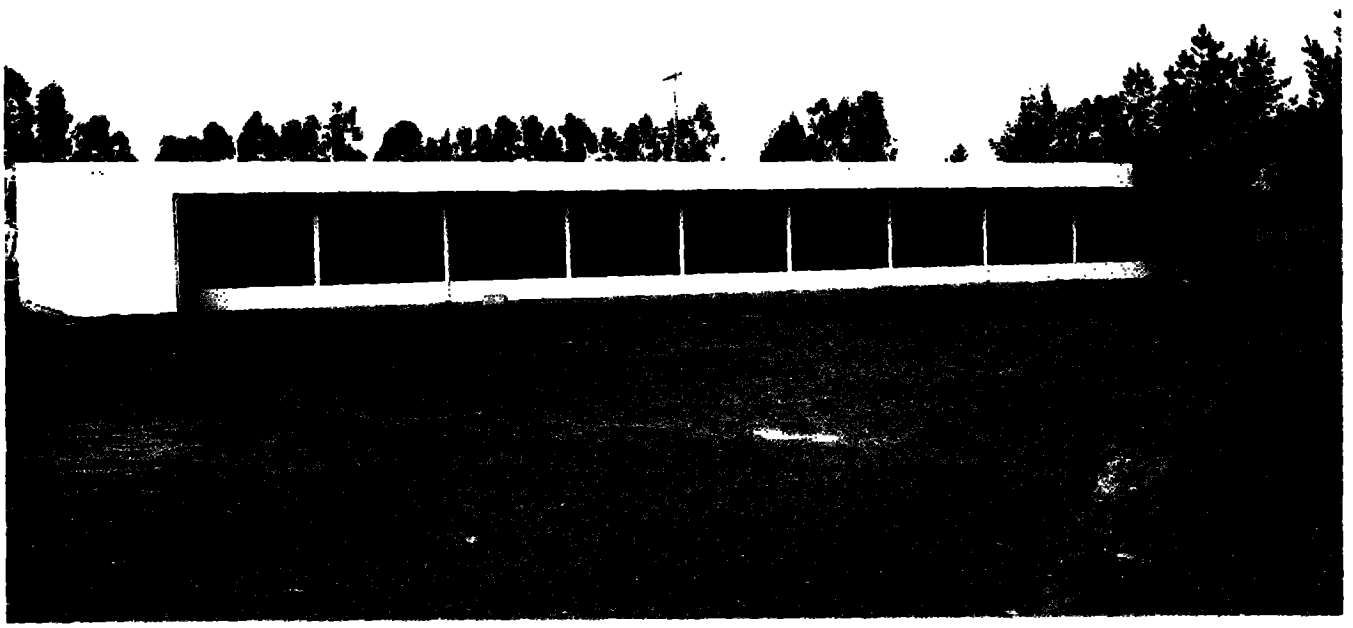

The application form

If the S-1 multiprocessor facility looks like the solution to your problems, please take a moment to tell us about your needs. Fill out this page or attach a letter and mail it to:

Technical Liaison Staff
S-1 Project
Lawrence Livermore National Laboratory
P.O. Box 808 L-276
Livermore, CA 94550

1. My computation problem is: _____
2. I need the following configuration of hardware:
   a. Uniprocessor or multiprocessor?
   b. How much disk space?
   c. Interactive graphics stations (resolution, color, speed characteristics)?
   d. Hardcopy graphics (resolution, color, speed characteristics)?
   e. How much magnetic tape storage?
   f. Special peripherals?
   g. Access by network or at Livermore?
   h. Security requirements?
3. I have the following questions or concerns about the S-1 multiprocessor facility: _____
4. My schedule is:
   a. How many hours of terminal connect time per week?
   b. When (please give preferred and alternate dates)?
5. My name is: _____
6. A second contact person is: _____

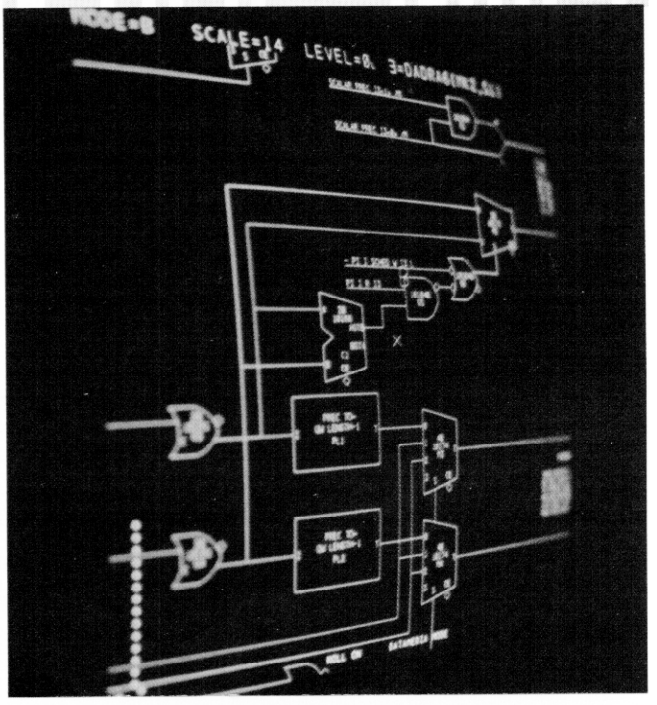

Logic drawings like this one, generated interactively on a computer terminal, serve as input to the CAD/CAM tools which made the design of the Mark IIA feasible.

Prepared for The Research and Technology Group, The Naval Electronic Systems Command, The U.S. Navy

Lawrence Livermore National Laboratory

LLL-TB-054